

Historical and future learning about climate sensitivity

Nathan M. Urban^{1,2}, Philip B. Holden³, Neil R. Edwards³, Ryan L. Sriver⁴, and Klaus Keller^{5,6}

Equilibrium climate sensitivity measures the long-term response of surface temperature to changes in atmospheric CO₂. The range of climate sensitivities in the IPCC AR5 Report is unchanged from that published almost 30 years earlier in the Charney Report. We conduct perfect-model experiments using an energy balance model to study the rate at which uncertainties might be reduced by observation of global temperature and ocean heat uptake. We find that a climate sensitivity of 1.5 °C can be statistically distinguished from 3 °C by 2030; 3 °C from 4.5 °C by 2040; and 4.5 °C from 6 °C by 2065. Learning rates are slowest in the scenarios of greatest concern (high sensitivities), due to a longer ocean response time, which may have bearing on wait-and-see vs. precautionary mitigation policies. Learning rates are optimistic in presuming the availability of whole-ocean heat data, but pessimistic by using simple aggregated metrics and model physics.

1. Introduction

A key climate uncertainty is the expected warming of the climate system with respect to a given change in atmospheric CO₂. This is often quantified by the long-term response of global surface temperature to an assumed doubling of CO₂ concentration, which is known as the (equilibrium) climate sensitivity [Knutti and Hegerl, 2008]. The climate sensitivity is uncertain not primarily because of scientific ignorance of the greenhouse effect, but because the climate system contains many complex dynamical feedbacks which modify the greenhouse warming. Greenhouse-induced changes in atmospheric water vapor content, cloud cover, snow and ice area, etc. themselves can cause warming or cooling effects which are absorbed into the definition of climate sensitivity.

Together, uncertainty about the strengths (or even signs) of these feedbacks lead to substantial uncertainty in the overall climate sensitivity, which was judged by the 2013 IPCC Fifth Assessment Report (AR4) to likely lie in the range of 1.5–4.5 °C per doubling of CO₂ [IPCC, 2013]. This is identical to the range reported in the influential 1979

Charney Report [Charney, 1979]. While the scientific understanding of the climate system has improved over time, this does not necessarily lead to reduced uncertainty about specific system properties such as climate sensitivity, as new complexities are discovered e.g. in the radiative forcing of aerosols. In addition, learning about the climate by temperature observations may be slow as the greenhouse effect has only gradually strengthened, and the oceans respond slowly to surface temperature perturbations.

This raises the question of when substantial reductions in uncertainty about climate sensitivity may be expected, i.e., the learning rate. The question is of potential policy interest: if learning rates are expected to be fast, then one can adopt a ‘wait-and-see’ approach, deferring strong emissions mitigation until more is learned. On the other hand, if learning rates are expected to be slow, then a ‘precautionary’ approach may be more advisable, with more rapid policy action taken to insure against the risk of a high climate sensitivity that cannot be quickly ruled out.

Quantitative studies of climate sensitivity learning have existed for some time, often reporting long learning times, though the definition of ‘learning’ as well as the methods used vary across studies. Rather than attempting to reconcile these different studies with each other and with our results, we simply summarize each approach and findings, highlighting where definitions and methods differ.

One early study, Khesghi and White [1993], fit an upwelling-diffusive energy-balance model (EBM) with autoregressive and white noise to the historical surface temperature record using least-squares minimization. It finds fairly tight uncertainty bounds ($\sim 1^\circ\text{C}$ 90% interval width) using only temperature data only up to 1980 and a 1.5-year autocorrelation time of natural variability (defined as the e-folding time of an exponential correlation function); however, if 100-year temperature variability is assumed, this level of certainty cannot be obtained until at least 2080 under the assumptions of the study.

Some subsequent work on climate learning arose from the economic integrated assessment modeling and policy community, using simpler energy balance models. A seminal study in endogenous Bayesian learning, Kelly and Kolstad [1999], estimated climate sensitivity sequentially using a Kalman-filter type approach, assuming normality in the climate feedback factor. It used an EBM with a single atmospheric layer and ocean layer — simpler than a diffusive ocean model — with simulated observations of both atmosphere and deep ocean temperature, including autoregressive temperature variability. Over many hypothetical realizations of future temperature, it found a learning time of 90–160 years, but used a rather strict requirement for ‘learning’ (a precision of 0.1° at 90% confidence). A more recent economic study, Leach [2007], considered joint learning of climate sensitivity and the autocorrelation of temperature variability using an EBM with atmospheric and ocean layers, similar to Kelly and Kolstad [1999]. It reports learning times, also with a precision of 0.1° , of hundreds to thousands of years depending on the emissions scenario.

More recent work in the climate science community has returned to the climate sensitivity learning problem. Padilla

¹Computational Physics and Methods (CCS-2), Los Alamos National Laboratory

²Woodrow Wilson School of Public and International Affairs, Princeton University

³Department of Environment, Earth and Ecosystems, Open University

⁴Department of Atmospheric Sciences, University of Illinois at Urbana-Champaign

⁵Department of Geosciences, Pennsylvania State University

⁶Department of Engineering and Public Policy, Carnegie Mellon University

et al. [2011] estimated the joint uncertainty in the transient climate sensitivity and aerosol forcing over the historical and near-future period using a two-layer EBM and a nonlinear Kalman filter, considering only a single assumed sensitivity in its future learning projections. It constrains these parameters using global surface temperature alone and neglects ocean heat uptake, absorbing uptake uncertainty into the definition of the transient feedback, arguing that this is less relevant to the estimation of the transient (as opposed to equilibrium) climate sensitivity. It finds that the lower bound on transient sensitivity remains relatively unchanged over time, but substantial reductions in the upper bound occur by 1950, and further rapid reductions occur after the year 2000 as recent temperature data exclude higher sensitivities. The implications for the upper bound of equilibrium sensitivity are not studied in this work.

Ring and Schlesinger [2007] produced an estimate of joint climate parameter learning (climate sensitivity, ocean diffusivity, and aerosol forcing) using an EBM with upwelling-diffusive ocean. It considers Bayesian learning from surface temperature and radiative imbalance from ocean heat uptake, using synthetic historical and future observations generated from the EBM and a singular spectrum analysis (SSA) of historical climate variability. It finds rapid learning about climate sensitivity (essentially perfect knowledge by 2050). However, its present-day uncertainty about climate sensitivity is already unusually narrow — approximately 0.5°C 90% interval width, compared to the 2.5°C 66% interval width of the IPCC.

Hannart et al. [2013] considers learning in a single-layer stochastic EBM along the lines of *Kelly and Kolstad* [1999], constrained by surface temperature data, and finds the possibility of ‘disconcerting learning’ where the uncertainty may (temporarily) increase over time as new observations are included. This effect is strongest when the prior distribution of climate uncertainty is highly skewed or heavy-tailed.




Figure 1. Historical observations and projected scenarios under four assumed climate sensitivities (S), for surface temperature (upper panel) and ocean heat (lower panel). At early times when observations are not available, synthetic hindcasts are generated using the same procedure used to generate the synthetic projections. Autoregressive noise is added to the model projections to simulate natural variability, observation error, and model structural error, as described in Sec. 2.1.

In this work, we conduct a climate sensitivity learning study with the following features: (1) a diffusive-ocean EBM; (2) joint Bayesian learning about both climate physics and stochastic climate noise (‘statistical’) parameters, with the latter determined endogenously from the data-model residuals; (3) assimilation of real historical temperature and ocean heat observations up to present, followed by synthetic observations in the future; (4) correlation of climate parameters in the synthetic future scenarios to enforce historical consistency; (5) dependence of learning rates upon the assumed climate sensitivity; and (6) a new learning criterion. In particular, we consider learning climate sensitivity to a tenth of a degree, as required in some earlier studies, to be too stringent a criterion for many decision-making contexts. Rather, we consider a relaxed learning criterion, which is the time to distinguish between a few coarse climate sensitivity scenarios (separated by 1.5°C intervals). Some of the features of our analysis have been considered individually in earlier work, but to date no work has combined all of them comprehensively.

2. Methods

We use temperature and ocean heat observations up to a given year to calibrate the parameters of an energy balance climate model (Fig. 1). The Bayesian calibration gives a probability distribution for climate sensitivity and other model parameters, conditional on the observations that exist up to the given year. To estimate the rate of learning, the calibration procedure is repeated by including additional observations into the analysis, 5 years at a time. The resulting sequence of parameter estimates gives the rate of learning about the model parameters over time as more data are included in the inference. The rate of learning for our purposes is defined to be reduction in width of the parameters’ 95% credible intervals over time.

2.1. Model and data sources

The underlying energy balance model (EBM), called DOECLIM, has a zero-dimensional atmosphere coupled to a one-dimensional diffusive ocean [*Kriegler, 2005; Tanaka et al., 2007*]. Its three uncertain parameters are climate sensitivity (S); ocean vertical diffusivity (κ), controlling the rate of heat ocean uptake and the climate response time; and the aerosol radiative forcing strength (α), expressed as a factor multiplying the historical or projected forcing time series. These three parameters are generally assumed in climate sensitivity studies to be the most dominant physical uncertainties in the global temperature response [*Forest et al., 2002; Knutti et al., 2002; Urban and Keller, 2010; Ring and Schlesinger, 2007*]. In addition, the initial temperature (T_0) and ocean heat anomalies (H_0) at the beginning of the model integration (the year 1850), and the standard deviation (σ) and annual autocorrelation (ρ) of the data-model residuals, are also treated as uncertain parameters to be estimated. Including both temperature and ocean heat observations, a total of 9 model and statistical parameters are jointly estimated: $\theta = [S, \kappa, \alpha, T_0, H_0, \rho_T, \rho_H, \sigma_T, \sigma_H]$.

The joint uncertainty in model and statistical parameters is estimated every 5 years from 1850 to 2010 by assimilating whatever temperature and ocean heat observations exist up to a given year. The posterior parameter probability distribution using data up to year t is denoted $[\theta|y_t]$, where θ is the vector of estimated parameters and y is the observational time series data for temperature and ocean heat. The surface temperature data are HadCRUT4 global means (1850–2012) [*Morice et al., 2012*], and the ocean heat data

are 0–3000 m ocean heat uptake (1953–1996) [Gouretski and Koltermann, 2007]. Although longer ocean heat time series are available [Levitus, 2012], they are for a depth range (0–2000 m) that is not compatible with the output of the DOE-CLIM model, which has a 4000 m ocean. In this analysis, we treat the observed 3000-m and modeled 4000-m heat uptake as comparable, assuming that little heat is transported to depths below 3000 m over the time scales considered. The DOECLIM model is forced with RCP 8.5 historical radiative forcings, with the total (direct+indirect) aerosol forcing multiplied by an uncertain factor estimated from the data, as mentioned above.

2.2. Parameter inference

The joint uncertainty in parameters is determined by Bayesian inference [Urban and Keller, 2010], so that the posterior distribution of the parameters conditional on the data is proportional to the likelihood function times the prior distribution, $[\theta|y_t] \propto [y_t|\theta][\theta]$. The Bayesian parameter posterior distribution is sampled by Markov chain Monte Carlo (MCMC) over 1 million iterations. The likelihood function $[y_t|\theta]$ assumes that the temperature and ocean heat observations are given by a deterministic mean trend simulated by the climate model, plus a normally-distributed first-order autoregressive noise process, i.e. AR(1), and the temperature and ocean heat residual noise processes are uncorrelated with each other, as in Urban and Keller [2010]. This residual AR(1) process represents a combination of both observation error, interannual natural variability, and model error. We do not attempt to disentangle and attribute these individual sources of variability, but rather assume that their sum can be modeled as an autoregressive process that can be estimated from the data-model misfit.

The temperature observations up to year t are thus given as $T_t = T_{EBM}(t; \theta) + \epsilon_T(\theta) + T_0(\theta)$ where T_{EBM} is the climate model prediction and ϵ_T is an AR(1) process; and similarly for the ocean heat observations.

Informative priors are used for the three EBM parameters, while uniform priors are used for the remaining six parameters; assuming prior independence, the joint prior $[\theta]$ is the product of these individual prior distributions. The prior probability distribution for climate sensitivity is loosely based on model climatological and Last Glacial Maximum constraints, and is equal to the product of two normal inverse Gaussian distributions [Olson et al., 2012], $NIG(\alpha = 1.8, \delta = 2.3, \beta = 1.2, \mu = 1.7) \times NIG(\alpha = 1.9, \delta = 3.3, \beta = 1.0, \mu = 1.3)$. The prior for vertical diffusivity is motivated by the range of values found in Goes et al. [2010]; Bhat et al. [2012], which constrained that parameter in the UVic intermediate complexity model using biogeochemical ocean tracer data. The observationally-permitted range of diffusivities found for UVic is converted to a range of effective diffusivities for the DOECLIM model by fitting DOECLIM to the transient ocean response of UVic; a Lognormal(1.1, 0.3²) distribution approximately spans this range. The prior for aerosol scaling is triangular over the range 0 to 3, peaked at 1, to conservatively represent the ranges summarized in the fourth IPCC report [IPCC, 2007]. These climatological and paleoclimatological constraints on climate sensitivity, ocean biogeochemical constraints on vertical diffusivity, and optical-depth based estimates of aerosol forcing are all approximately independent of the historical surface temperature and ocean heat data sets and can serve as useful prior knowledge about the EBM parameters.

The priors chosen here are not intended to reflect historical scientific knowledge about the climate system (i.e., in the mid-19th century). The ‘learning’ estimated over the period spanned by the instrumental record should not be interpreted to summarize the evolution of scientific understanding of climate sensitivity. Rather, it is intended to demonstrate how the accumulation of temperature and ocean heat data over time can constrain climate parameters, relative

to some independent state of knowledge represented by one particular informed choice of prior.

2.3. Future learning

To analyze rates of future learning, the MCMC-based sequential inference is repeated into the future with synthetic (model-generated) observational data. Four future scenarios are considered, with low (1.5 °C), medium (3 °C), high (4.5 °C), and very high (6 °C) climate sensitivities (Fig. 1).

The synthetic observations for each scenario are generated by integrating the energy balance model forward with RCP 8.5 extended forcings to 2150, and adding simulated AR(1) noise to the model projections. A high emissions scenario is chosen to represent the case that society does not choose to immediately implement strict carbon controls, which more strongly brings out the importance of learning to the wait-and-see vs. precautionary policy contrast.

Each of the four scenarios assumes a particular climate sensitivity. Because the observational constraints imply correlations between climate sensitivity and other parameters, it is necessary to set the other parameters in each scenario to values that are consistent with the assumed sensitivity. For example, a high-sensitivity scenario will generally require a high rate of ocean heat uptake or a strong aerosol cooling in order to produced the observed historical warming [Urban and Keller, 2009, 2010].

To generate consistent parameter sets, the parameters other than climate sensitivity are fixed at their posterior means given observations to the year 2012 (the year before the projections begin), conditional on the assumed climate sensitivity, Fig. 2. That is, denoting climate sensitivity by S and the remaining 8 parameters by ϕ so that $\theta = (S, \phi)$, an estimate for ϕ conditional on S given data up to year $t = 2012$ is $\bar{\phi}_{2012|S} = E[\phi|y_t, S]$. Geometrically the conditional distribution $[\phi|y_t, S]$ can be thought of as the slice of the joint posterior distribution $[\phi|y_t]$ generated by fixing S

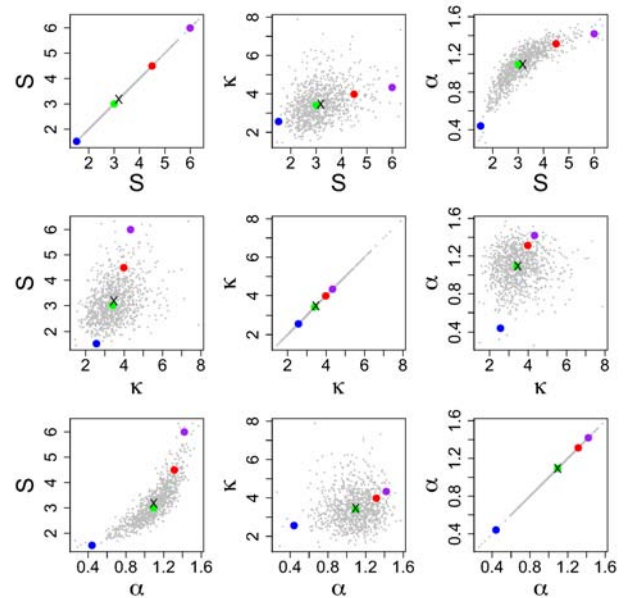


Figure 2. Pairwise correlations between the climate sensitivity (S), vertical diffusivity (κ), and aerosol scaling (α) estimated in 2012. Gray points are samples from the posterior distribution; crosses indicate the posterior mean; colored circles indicate the conditional mean values used in the four future scenarios in Fig. 1.

to a constant. The conditional mean is approximated numerically by taking a thin slice of the Monte Carlo samples from the joint posterior, $[\theta|y_{2012}]$, that are within $\pm 0.1^\circ\text{C}$ of the assumed climate sensitivity S , and calculating the average of ϕ over this subset. This procedure gives values of the vertical diffusivity and aerosol forcing, as well as the statistical parameters, that are consistent with both the assumed climate sensitivity and observational data.

3. Results

Inferred marginal posterior probability distributions for the three climate parameters given data to 2012 are shown in Fig. 3. The historical temperature and ocean heat data constrain the climate sensitivity and ocean diffusivity parameters only minimally relative to their assumed prior distributions. The prior and posterior distributions for climate sensitivity are similar to the $2\text{--}4.5^\circ\text{C}$ range given in the IPCC Fourth Assessment Report [IPCC, 2007], although it should be noted that Fig. 3 depicts 90/95% ranges while the IPCC report implicitly uses a 66% range; our 66% range for 2012 is $2.5\text{--}3.9^\circ\text{C}$. The posterior aerosol forcing scaling factor does show a substantial decrease in uncertainty after updating with observational data, indicating that the historical warming response of the global climate system indirectly constrains the possible aerosol forcing.

The learning results are summarized in Fig. 4. As implied by Fig. 3, little learning is observed to the present date (relative to the assumed priors) for climate sensitivity and ocean diffusivity. We hypothesize this is due to historical cancellation between the greenhouse and aerosol forcings, although some learning occurs about the aerosol forcing itself.

Into the future, substantial learning occurs over time about all three climate parameters, for all four projected climate sensitivity scenarios. The uncertainty in climate sensitivity, as measured by the width of the 95% credible interval, is reduced from 3.1°C in 2012 to $0.4\text{--}2.6^\circ$ in 2050,




Figure 3. Posterior distributions of the three estimated climate parameters given data to 2012. The solid black curves are the posteriors, the green dashed curves are the priors, the solid circles are the corresponding means, and the thick/thin line segments are the 90/95% intervals.

$0.1\text{--}0.5^\circ$ in 2100, and $0.1\text{--}0.2^\circ$ in 2150, depending on the assumed climate sensitivity. (A precision of 0.1° may be unrealistic given the highly simplified assumptions of the climate physics and model structural error.) Reductions in uncertainty also occur for ocean diffusivity and aerosol forcing, although learning about aerosol forcing ceases as the forcing itself diminishes in the RCP 8.5 scenario.

In general, more rapid learning about climate sensitivity occurs for lower sensitivities. This occurs because lower sensitivities are also associated with faster response times (holding vertical diffusivity fixed), as derived in Hansen *et al.* [1985]. This can also be seen in Fig. 1, as the warming projected in the high-sensitivity scenarios are more similar to each other than in the low-sensitivity scenarios, implying that it is more difficult to learn about high climate sensitivities than low sensitivities. Also, we might expect slower learning about unlikely extreme scenarios in the tails of the present-day probability distribution, such as a very high climate sensitivity, under the expectation that more data may be required to contradict our beliefs about scenarios we judge to be unlikely than ones we judge to be likely. However, the low climate sensitivity scenario is also in the tail of the present-day distribution, and experiences faster learning rates than the very high scenario.

To explore the robustness of slow learning about high climate sensitivities with respect to model formulation, we compare the transient surface warming projected for 2100 as a function of climate sensitivity across four different models

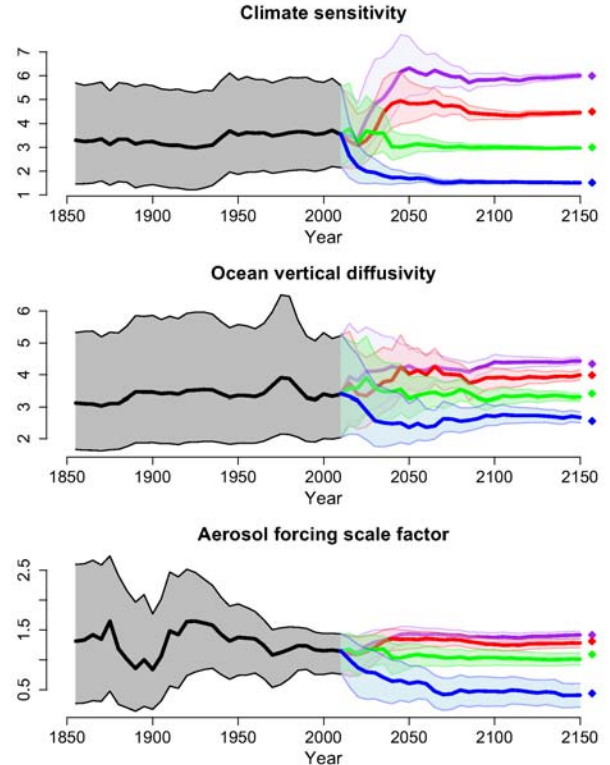


Figure 4. Historical and future learning over time about the three estimated climate parameters. Black indicates historically constrained estimates; the four colors represent the four future sensitivity scenarios in Fig. 1. The thick curves are the posterior means, and the shaded bands are the marginal 95% intervals. The diamonds on the right are the assumed ‘true’ values in the four future learning scenarios.

(Fig. 5): a simple 2-layer EBM from the DICE-2007 integrated assessment model [Nordhaus, 2007] (adapted from Schneider and Thompson [1981]), the DOECLIM diffusive-ocean EBM used in this study, and the GENIE model of intermediate complexity [Holden et al., 2010].

Both the 2-layer EBM and the GENIE model agree quite closely with the predictions of the DOECLIM model. All the curves exhibit the feature that high sensitivities have similar transient responses, leading to slower learning. In all but the black dashed curve, all parameters but climate sensitivity are held fixed, which simplifies the analysis but does not correspond to the methodology of this paper. By contrast, our analysis allows the other parameters to be correlated with climate sensitivity as in Fig. 2, as implied by historical observational constraints.

To study how observational constraints on non-climate sensitivity parameters affect the transient temperature response, DOECLIM experiments are carried out that allow the ocean diffusivity and aerosol forcing parameters to covary with climate sensitivity as described in Section 2.3. In general, this reduces the expected warming (black dashed curve) compared to holding ocean diffusivity and aerosol forcing constant at their 2012 posterior mean values. (The noisy behavior of the curve at high climate sensitivities occurs because there are few posterior samples in 2012 with high sensitivities from which to compute the mean transient response.) However, the qualitative feature remains that high sensitivities are more difficult to distinguish from each other than are low sensitivities on the basis of the observed temperature response.

4. Conclusions

Under the assumptions of this study, substantial learning about climate sensitivity is possible during this century,




Figure 5. Transient warming (as measured by surface temperature difference between 1850 and 2100) as a function of equilibrium climate sensitivity, for a 2-layer energy balance model (DICE-2007), the diffusive ocean energy balance model used in this study (DOECLIM), and an Earth system model of intermediate complexity (GENIE). For DOECLIM, two cases are considered: where the ocean diffusivity and aerosol forcing are held fixed, and where they are allowed to covary with climate sensitivity consistent with historical observations.

reducing the 95% interval width to 0.5° or less. Climate sensitivity scenarios, separated into coarse 1.5° intervals from 1.5 – 6°C can be distinguished from each other with ~ 20 – 55 additional years of observations, depending on scenario. The climate sensitivity scenarios can be statistically distinguished from the next highest scenario, insofar as their 95% credible intervals no longer overlap, by 2030 to 2065, depending on scenario.

However, these learning rates are dependent on the assumption that the ocean heat uptake is observable for the entire ocean. Most data sets typically only give data for the upper 700 or 2000 meters (e.g., Levitus [2012], with the




Figure 6. Projected temperature and ocean heat, replicating Fig. 1 for RCP 8.5 (left), compared to RCP 2.6 (right).

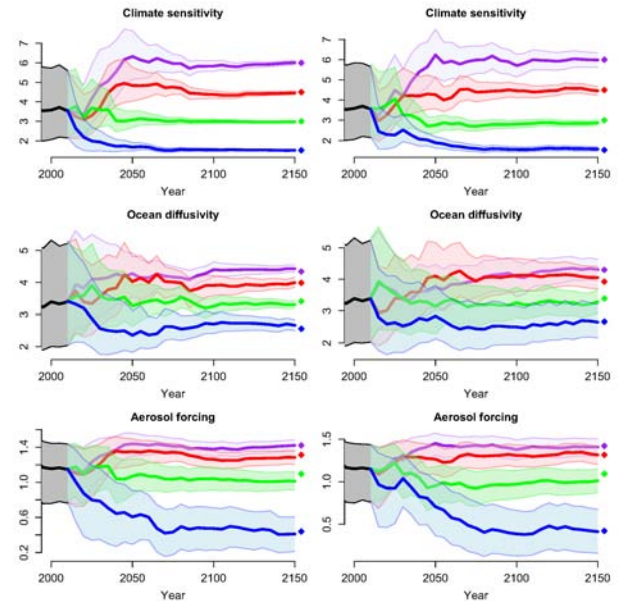


Figure 7. Parameter learning, replicating Fig. 4 for RCP 8.5, compared to RCP 2.6 (right).

accuracy of the latter decreasing significantly in years prior to the last decade. Learning rates should be slower with more realistic assumptions about the limitations of ocean heat observations. Learning will also be slower under lower emissions scenarios than the assumed RCP 8.5, although the time to distinguish between climate sensitivities is not greatly reduced, which we hypothesize is due to the large ocean heat signal-to-noise ratio present across RCP scenarios (Appendix A). This analysis makes strong simplifying assumptions about climate physics (by using a global EBM) as well as the nature of the model structural error (by using an AR(1) process fit to the data-model residuals). These assumptions could also lead to optimistic uncertainty estimates.

On the other hand, the analysis here uses only a simple EBM with globally- and annually-averaged surface temperature and ocean heat constraints. It neglects all spatial information [Sansó and Forest, 2009], the seasonal cycle [Knutti et al., 2006], shortwave and longwave radiation observations [Huang et al., 2010; Feldman et al., 2011; Huber et al., 2011], dynamical constraints on feedbacks [Fasullo and Trenberth, 2012], and other observational constraints that could be brought to bear on the problem. Nor does it analyze feedbacks at the regional level or decomposes climate sensitivity into individual component feedback processes (clouds, water vapor, surface albedo, etc.).

Thus, in the absence of a more sophisticated analysis, the question of whether the learning rates in this paper are optimistic or pessimistic remains open. The present analysis also neglects uncertainty about the learning rate itself; conditioning on the expected model parameters and only sampling a single projected future for each gives only a noisy point estimate of the learning time rather than a distribution of possible learning times; such distributions have been considered in earlier work, such as Kelly and Kolstad [1999]. Followup work should also revisit the policy implications found in earlier integrated economic analyses, to determine how the learning results found here may shift the balance between immediate vs. deferred mitigation.

Acknowledgments. We would like to acknowledge Roman Olson for helpful discussions concerning the choice of priors. This

work was supported under Princeton University's Program in Science, Technology and Environmental Policy; the Office of Science (BER), U.S. Department of Energy; and the National Science Foundation through the Network for Sustainable Climate Risk Management (SCRiM) under NSF cooperative agreement GEO-1240507. Data used in this paper are available from the corresponding author upon request.

Appendix A: Other learning assumptions

A less strongly forced system will lead to a weaker climate response, which in principle should result in longer learning times. To test this hypothesis, we repeat the learning study using the low RCP 2.6 emissions scenario. The projections, compared to the RCP 8.5 case in the main text, are shown in Fig. 6, and the parameter learning in Fig. 7.

Although there are some differences in climate sensitivity learning rate across the RCP 8.5 and 2.6 scenarios, the time required to differentiate between different sensitivities is qualitatively similar under both high and low emissions projections. We hypothesize this is due to the large signal-to-noise ratio present in the simulated ocean heat uptake projections. Intuitively, if we expect that the ocean heat data provide a strong constraint on climate sensitivity, the learning time will be roughly determined by the time it takes for the projected ocean heat curves to no longer overlap with each other. The projections in the lower panels of Fig. 6 suggest the uptake curves may be discriminated around ~ 2050 , consistent with the parameter learning results in Fig. 7.

To highlight the importance of the ocean heat constraint, Fig. 8 shows learning results that ignore the ocean heat data, only constraining the parameters with surface temperature. Substantially delayed learning results, with the 3°C scenario separating from higher sensitivities only after 2100, and the two higher scenarios distinguishable from each other around 2150. The ocean diffusivity parameter, as expected, not practically identifiable using only surface temperature data (little learning occurs).

References

- Bhat, K. S., M. Haran, R. Olson, and K. Keller (2012), Inferring likelihoods and climate system characteristics from climate models and multiple tracers, *Environmetrics*, 23, 345–362, doi:10.1002/env.2149.
- Charney, J. G. e. a. (1979), Carbon dioxide and climate: A scientific assessment, *Tech. rep.*, National Research Council, Woods Hole, MA, USA.
- Fasullo, J. T., and K. E. Trenberth (2012), A less cloudy future: The role of subtropical subsidence in climate sensitivity, *Science*, 338, 792–794, doi:10.1126/science.1227465.
- Feldman, D. R., C. A. Algieri, W. D. Collins, Y. L. Roberts, and P. A. Pilewskie (2011), Simulation studies for the detection of changes in broadband albedo and shortwave nadir reflectance spectra under a climate change scenario, *J. Geophys. Res.*, 116, D24103, doi:10.1029/2011JD016407.
- Forest, C. E., P. H. Stone, A. P. Sokolov, M. R. Allen, and M. D. Webster (2002), Quantifying uncertainties in climate system properties with the use of recent climate observations, *Science*, 295, 113–117, doi:10.1126/science.1064419.
- Goes, M., N. M. Urban, R. Tonkonojnikov, M. Haran, and K. Keller (2010), What is the skill of ocean tracers in reducing uncertainties about ocean diapycnal mixing and projections of the Atlantic Meridional Overturning Circulation?, *J. Geophys. Res.*, 115, C12006, doi:10.1029/2010JC006407.
- Gouretski, V., and K. P. Koltermann (2007), How much is the ocean really warming?, *Geophys. Res. Lett.*, 34, L01610, doi: 10.1029/2006GL027834.
- Hannart, A., M. Ghil, J.-L. Dufresne, and P. Naveau (2013), Disconcerting learning on climate sensitivity and the uncertain future of uncertainty, *Climatic Change*, 119, 585, doi: 10.1007/s10584-013-0770-z.

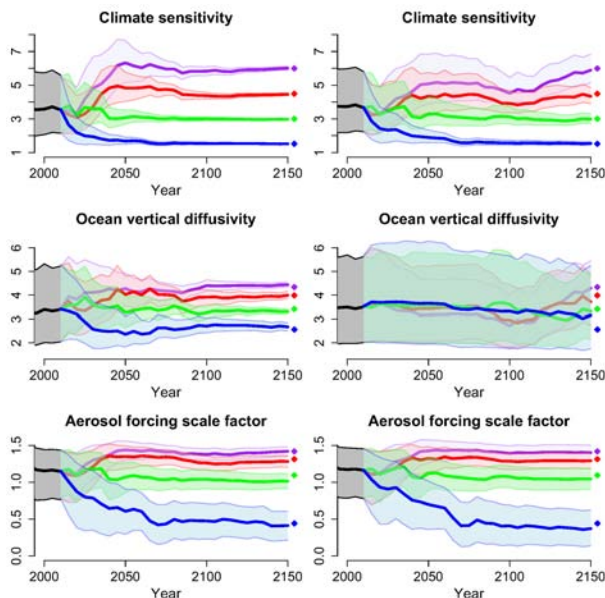


Figure 8. Parameter learning, replicating Fig. 4 for surface temperature and ocean heat constraints (left), compared to only surface temperature as a constraint (right).

- Hansen, J., G. Russell, A. Lacis, I. Fung, D. Rind, and P. Stone (1985), Climate response times: Dependence on climate sensitivity and ocean mixing, *Science*, *229*, 857–859, doi:10.1126/science.229.4716.857.
- Holden, P. B., N. R. Edwards, K. I. C. Oliver, T. M. Lenton, and R. D. Wilkinson (2010), A probabilistic calibration of climate sensitivity and terrestrial carbon change in GENIE-1, *Clim. Dyn.*, *35*, 785806, doi:10.1007/s00382-009-0630-8, 2010.
- Huang, Y., S. Leroy, P. J. Gero, J. Dykema, and J. Anderson (2010), Determining longwave forcing and feedback using infrared spectra and GNSS radio occultation, *J. Climate*, *23*, 6027–6035, doi:10.1175/2010JCLI3588.1.
- Huber, M., I. Mahlstein, and M. Wild (2011), Constraints on climate sensitivity from radiation patterns in climate models, *J. Climate*, *24*, 1034–1052, doi:10.1175/2010JCLI3403.1.
- Kelly, D. L., and C. D. Kolstad (1999), Bayesian learning, pollution, and growth, *J. Economic Dynamics and Control*, *23*, 491–518, doi:10.1016/S0165-1889(98)00034-7.
- Khesghi, H. S., and B. S. White (1993), Does recent global warming suggest an enhanced greenhouse effect?, *Climatic Change*, *23*, 121–139, doi:10.1007/BF01097333.
- Knutti, R., and G. C. Hegerl (2008), The equilibrium sensitivity of the Earth’s temperature to radiation changes, *Nat. Geosci.*, *1*, 735–743, doi:10.1038/ngeo337.
- Knutti, R., T. F. Stocker, F. Joos, and G. K. Plattner (2002), Constraints on radiative forcing and future climate change from observations and climate model ensembles, *Nature*, *416*, 719–723, doi:10.1038/416719a.
- Knutti, R., G. A. Meehl, M. R. Allen, and D. A. Stainforth (2006), Constraining climate sensitivity from the seasonal cycle in surface temperature, *J. Climate*, *19*, 4224–4233, doi:10.1175/JCLI3865.1.
- Kriegler, E. (2005), Imprecise probability analysis for integrated assessment of climate change, *Tech. rep.*, University of Potsdam, ph.D. thesis.
- Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K. B. Averyt, M. Tignor, and H. L. Miller (Eds.) (2007), Summary for Policymakers. In: *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Stocker, T. F., D. Qin, G.-K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. M. Midgley (Eds.) (2013), Summary for Policymakers. In: *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chap. Summary for Policymakers, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Leach, A. J. (2007), The climate change learning curve, *J. Economic Dynamics and Control*, *31*, 1728–1752, doi:10.1016/j.jedc.2006.06.001.
- Levitus, S. e. a. (2012), World ocean heat content and thermoclinic sea level change (0–2000 m), 1955–2010, *Geophys. Res. Lett.*, *39*, L10603, doi:10.1029/2012GL051106.
- Morice, C. P., J. J. Kennedy, N. A. Rayner, and P. D. Jones (2012), Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: the HadCRUT4 dataset, *J. Geophys. Res.*, *117*, D08101, doi:10.1029/2011JD017187.
- Nordhaus, W. D. (2007), *A Question of Balance: Weighing the Options on Global Warming Policies*, Yale University Press, New Haven, CT, USA.
- Olson, R., R. Sriver, M. Goes, N. M. Urban, H. D. Matthews, M. Haran, and K. Keller (2012), A climate sensitivity estimate using Bayesian fusion of instrumental observations and an Earth System model, *J. Geophys. Res.*, *117*, D04103, doi:10.1029/2011JD016620.
- Padilla, L. E., G. K. Vallis, and C. W. Rowley (2011), Probabilistic estimates of the transient climate sensitivity subject to uncertainty in forcing and natural variability, *J. Climate*, *24*, 5521, doi:http://dx.doi.org/10.1175/2011JCLI3989.1.
- Ring, M. J., and M. E. Schlesinger (2007), Bayesian learning of climate sensitivity i: Synthetic observations, *Atmos. Climate Sci.*, *2*, 464–473, doi:10.4236/acs.2012.24040.
- Sansó, B., and C. Forest (2009), Statistical calibration of climate system properties, *J. Roy. Stat. Soc. C*, *58*, 485–503, doi:10.1111/j.1467-9876.2009.00669.x.
- Schneider, S., and S. L. Thompson (1981), Atmospheric CO₂ and climate: Importance of the transient responses, *J. Geophys. Res.*, *86*, 3135–3147, doi:10.1029/JC086iC04p03135.
- Tanaka, K., E. Kriegler, T. Bruckner, G. Hooss, W. Knorr, and T. Raddatz (2007), Aggregated carbon cycle, atmospheric chemistry, and climate model (ACC2), *Tech. Rep. Reports on Earth System Science No. 40*, Max Planck Institute for Meteorology, Hamburg, Germany.
- Urban, N., and K. Keller (2009), Complementary observational constraints on climate sensitivity, *Geophys. Res. Lett.*, *36*, L04708, doi:10.1029/2008GL036457.
- Urban, N. M., and K. Keller (2010), Probabilistic hindcasts and projections of the coupled climate, carbon cycle and Atlantic meridional overturning circulation system: A Bayesian fusion of century-scale observations with a simple model, *Tellus A*, *62*, 737–750, doi:10.1111/j.1600-0870.2010.00471.x.

Corresponding author: Nathan Urban, Computational Physics and Methods (CCS-2), Los Alamos National Laboratory, P.O. Box 1663, MS B265, Los Alamos, NM 87545, USA. (nurban@lanl.gov)